APPROACHING THE GIGABIT NETWORKING THRESHOLD IN THE U.S.*

Steven N. Goldstein
National Science Foundation, USA

Abstract
When the Gigabit Network Testbed Project began in the United States in 1989, Megabit/second (Mbps) transmission rates were considered fast for wide-area networking, and people who spoke of Gigabits/second (Gbps) were considered visionary. Now, with the final report on the Testbed Project almost ready for publication, production-quality Gigabit networking is not far off in the United States. In my talk, I will review the key findings of the Gigabit Testbed Project, and then I will highlight several developments that are leading to production-level Gigabit networking. These include: commercial backbones running at OC3c (155 Mbps) and OC12c (622 Mbps); a White House initiative for a "Next Generation Internet" as well as a multi-university initiative for "Internet II" both of which would reach Gigabit speeds; a university-developed Gigabit/ATM test kit and an associated cooperative research opportunity; efforts to develop "central office grade" Internet switches capable of multi-Gbps operation, and ATM-over-satellite (up to155 Mbps), including Japan-U.S. cooperative efforts on high data rate satellite transmission experiments. I will also mention experiments with the CAVE, an immersive three-dimensional virtual reality environment that can operate in distributed mode over broadband networks. (Both the Gigabit/ATM Test Kit and the CAVE technologies can be obtained on a cost-recovery basis for intercontinental research collaborations.)

---

* Prepared for presentation at the *The International Symposium of Gigabit Networks,* organized by the Research Support Center for Advanced Telecommunications, Technology Research Foundation and held at Kokuyo Hall, Sinagawa-ku, Tokyo, Japan, January 29, 1997

APPROACHING THE GIGABIT NETWORKING THRESHOLD IN THE U.S.

This is an overview of United States research, development and commercial implementation of high performance networking that operates at or near gigabit per second (Gbps) transmission rates. The central message is that the gigabit threshold has been passed in the realm of research and development, and that pre-commercial and commercial implementations are poised to cross the gigabit threshold before the end of the millennium. What follows is largely a compilation of material either provided directly by investigators of high performance networking or obtained from open sources, often with their guidance. The author is indebted to colleagues who have cooperated to help make this information available to the "International Symposium of Gigabit Networking." The author assumes full responsibility, however, for any errors of omission and mis-interpretation.

GIGABIT NETWORK TESTBED PROJECT [1]

Introduction

The final Gigabit Network Testbed Initiative report is still in draft form. Therefore, with the kind permission and cooperation of the Corporation for National Research Initiatives, much of the information contained in the Executive Summary is reproduced here for the benefit of the International Symposium of Gigabit Networking.

The Gigabit Network Testbed Initiative was a major effort by approximately forty organizations representing universities, telecommunication carriers, industry and national research laboratories, and computer companies to create a set of very high-speed network testbeds and to explore their application. This effort, funded by the National Science Foundation (NSF) and the Defense Advanced Research Projects Agency (DARPA), was orchestrated by the Corporation for National Research Initiatives (CNRI) working closely with each of the participating organizations and the U.S. Government. The U.S. Government seeded this effort with a total of approximately $20M over a period of approximately five years, with these funds used primarily to fund the university research efforts. Major transmission facilities and equipment were donated at no cost to the project by the carriers and computer companies, who also directly funded their participating researchers and contributed to the university funding. The total value of industry contributions to the effort was estimated to be about $400M.

---

[1]This material is adapted from a draft of the Executive Summary of the forthcoming final report on the Gigabit Testbed Network by Robert Kahn and Richard Binder, Corporation for National Research Initiatives, Reston, Virginia, USA.

In addition to participants from leading universities, the participants included the Department of Energy's Lawrence Berkeley Laboratory, and Los Alamos National Laboratory (LANL); the National Astronautics and Space Administration's Jet Propulsion Laboratory; and the NSF-sponsored National Center for Supercomputer Applications, Pittsburgh Supercomputer Center, and San Diego Supercomputer Center. Industry research contributors included IBM Research, Bellcore, GTE Laboratories, AT&T Bell Laboratory, BellSouth Research, and MCNC.

Five Testbeds, named Aurora, Blanca, Casa, Nectar and Vistanet, were established and used between 1989 and 1995. The Testbeds were used to explore advanced networking issues and to investigate the alternatives for Gigabit network architectures in the future, and to carry out a wide range of experimental applications in areas such as weather modeling, chemical dynamics, radiation oncology, and telemetry modeling.

A summary of each of the Testbeds is given in Figures 1-5.

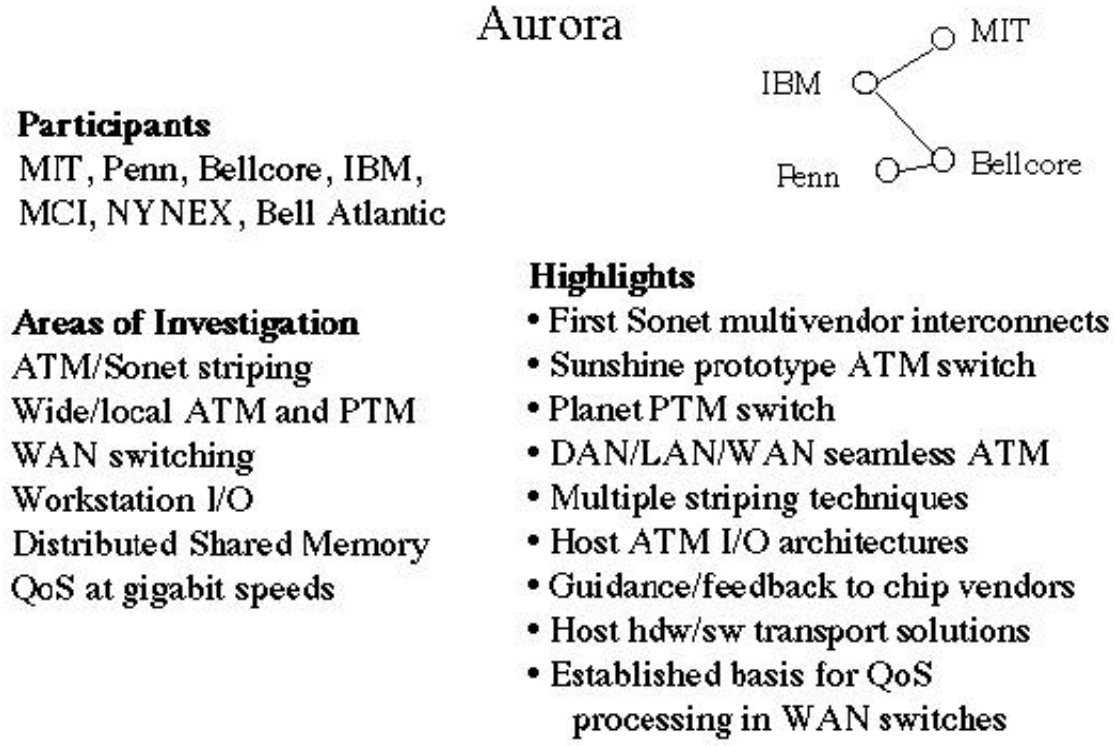

Figure 1. Aurora Testbed

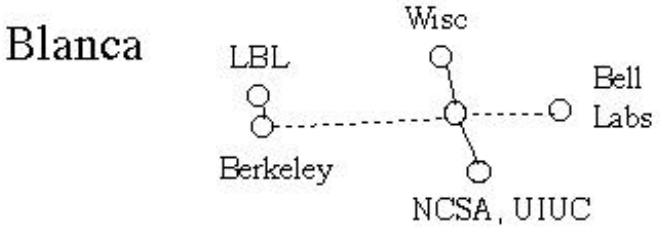

Participants
Wisconsin, NCSA, UIUC, Berkeley, LBL, AT&T

Areas of Investigation
Remote visualization and modeling
WAN congestion ctrl
Real-time protocols
ATM switch control
WAN/LAN interworking
Distributed shared memory

Highlights
- All-optical WAN transmission
- Tenet real-time protocol suite
- MPP I/O characterization
- 50x presentation layer speedup
- DSM latency-hiding techniques
- HIPPI-ATM gateway
- Dynamic time windows protocol
- Remote thunderstorm visualization
- Large-application I/O rates

Figure 2. Blanca Testbed

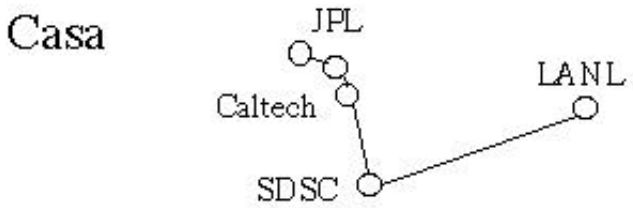

Participants
Caltech, JPL, LANL, SDSC, Parasoft, MCI, Pacific Bell

Areas of Investigation
Heterogeneous metacomputing
HIPPI-Sonet networking
Application support software
Network performance evaluation
Long-distance Sonet operation

Highlights
- Distributed application partitioning
- Superlinear computing speedup
- 1000-mile 2.5Gbps Sonet
- 1.2Gbps to endusers
- HIPPI-Sonet gateways with variable-bandwidth striping
- 500Mbps TCP, 790Mbps HIPPI over 1000-mile Sonet link
- Express appl. control software
- Outboard TCP/SHIP protocols

Figure 3. Casa Testbed

Nectar

Participants
CMU, Pittsburgh Supercomputer Center, Bellcore, Bell Atlantic

Areas of Investigation
Distr. application programming environments
WAN ATM/Sonet access
MPP distr. memory I/O
Workstation I/O
Distributed applications

Highlights
- Prototype WAN HIPPI-ATM-Sonet gwy
- Workstation 622Mbps TCP/IP requirements
- Optimal data distribution strategies for efficient distributed memory I/O
- Control software and BW vs Processing experiments for network of workstations
- Program development software for dynamic load-balancing and checkpointing
- Distribution of chemical plant processing application

Figure 4. Nectar Testbed

Vistanet

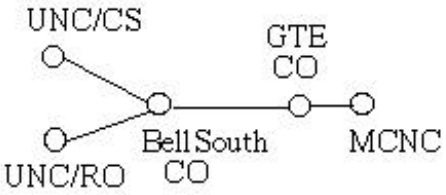

Participants
UNC, MCNC, GTE, BellSouth, NCSU

Areas of Investigation
Distributed interactive modeling
Graphics software for networked applications
Wide area ATM/Sonet
HIPPI-ATM-Sonet interworking
Transport protocol performance
Traffic measurement and characterization

Highlights
- Cancer treatment planning application demonstration
- OC-12c Sonet crossconnect
- First CO broadband ATM switch
- OC-12c links to user sites
- HIPPI-ATM/Sonet OC-12c adapter
- Outboard TCP/IP device
- 800 Mbps HIPPI traffic capture
- ATM application data capture and switch congestion modeling
- TCP vs XTP evaluation

Figure 5. Vistanet Testbed

Some of the key technical issues addressed in the program were transmission, switching, interworking, host I/O, network management, applications and support tools. In each case, various approaches were identified and many were considered in detail. A key decision in the program was to focus on technologies which were soon to be made available by industry or were at least working in the laboratory. ATM, SONET and HIPPI were three such technologies considered in the program.

Impacts

In addition to the many technical contributions resulting from the testbeds, a number of non-technical results have had major impacts for both education and industry.

First and foremost was the new model for network research provided by the testbed initiative. The bringing together of network and application researchers, integration of the computer science and telecommunications communities, academia-industry-government research teams, and government-leveraged industry funding, all part of a single, orchestrated project spanning the country, provided a new level of research collaboration not previously seen in this field. The Initiative created a high performance networking community that crossed academic/industry/government boundaries.

The coupling of application and networking technology research from project inception was a major step forward for both new technology development and applications progress. Having applications researchers involved from the start of the project allowed networking researchers to obtain early feedback on their network designs from a user's perspective, and allowed network performance to be evaluated using actual user traffic. Similarly, application researchers could learn how the network impacts their distributed application designs through early deployment of prototype software. Perhaps most significantly, they could proceed to investigate networked application concepts without first waiting for the new networks to become operational, opening their minds to new possibilities after decades of constrained bandwidth.

The coupling of packet networking researchers, who have largely come from the field of computer science, with the carrier telecommunications community provided another important dimension of integration. The development of computer communications networks and carrier-operated networks have historically proceeded along two separate paths with relatively little cross-fertilization. The testbeds allowed each community to work closely with the other, allowing each to better appreciate the other's problems and solutions and leading to new concepts of integrated networking and computing.

Another important dimension of the testbed model was its funding structure, in which government funding leveraged a large investment by industry. The major industry contribution was made by the carriers in the form of wide area SONET and other transmission facilities and equipment within each testbed operating at Gigabit or near-Gigabit rates. Not only were such services non-existent at the time, but they would have been unaffordable if they had existed. The carriers were willing to provide the facilities and to fund the participation of their researchers and engineers because of the opportunity to learn about potential applications of high speed networks while also benefiting from collaboration with the government-funded researchers in network technology experiments.

The Initiative also resulted in significant technology transfers to the commercial sector. As a direct result of their participation in the early stages of the project, two of the researchers at Carnegie-Mellon University founded a local-area ATM switch startup company, Fore Systems. This was the first such company formed, and provided a major stimulus for the emergence of high speed local area networking products. It also introduced to the product marketplace the integration of advanced networking concepts with advanced computing architectures used within their switch.

Other significant technology transfers involved the Hilda HIPPI measurement device developed as part of the Vistanet effort and the HIPPI-SONET wide area gateway developed by LANL for the Casa testbed. Both of these systems have been successfully commercialized by the private sector. In other cases, new high speed networking products were developed by industry in direct response to the needs of the testbeds, for example HIPPI fiber optic extenders and high speed user-side SONET equipment. Additionally, major technology transfers occurred through the migration of students who had worked in the testbeds to industry to implement their work in company products.

Investigations

At least four distinctive end-to-end network layer architectures were used as research platforms in the testbeds. This was a result both of architecture component choices made by researchers after the work was underway and of the a priori testbed selection process. Figure 6 gives a high-level view of these end-to-end architectures.

A. Seamless Wide-Local Area ATM

Figure 6-A, Seamless Wide-Local Area ATM, reflects one of the resulting architectures which was investigated in the Aurora testbed. In this case all networks use ATM cell switching and are interconnected without use of an IP layer. Sonet is used as the underlying transmission technology for the wide area network, and in some cases also for local distribution. In other cases other transmission technologies are used under ATM in the local area. In addition to the usual notion of LANs, Aurora introduced ATM Desk Area Networks, or DANs, as part of the overall architecture.

B. Heterogeneous Wide Area ATM and Local Area Technologies

Figure 6-B, Heterogeneous Wide Area ATM and Local Area technologies, shows the architecture used in the Blanca, Nectar, and Vistanet testbeds. ATM is again the switching technology used in the wide area network, but a non-ATM technology (in these instances, HIPPI) is used for local area connectivity. Both bridging and gateway approaches were investigated, and wide area transmission included SONET and all-optical infrastructures.

C. Seamless Wide-Local Area PTM

Figure 6-C, Seamless Wide-Local Area PTM, is the Packet Transfer Mode analog to the seamless ATM architecture and, like the latter, was part of the Aurora testbed work. In this case variable-length packets are forwarded across both the local and wide area networks, which were designed to operate as an integrated system.

D. Wide Area HIPPI via Local Switching

Figure 6-D, Wide Area HIPPI via Local Switching, reflects the resulting Casa testbed architecture. In this case HIPPI was used for both local and wide area switching, with SONET providing the wide area transmission infrastructure. The combination of one or more specially designed gateways at each site, in conjunction with a site's local HIPPI switch, provided the wide area routing/switching structure for variable-length packet forwarding through intermediate sites.

Selected Results

The full report should be read to capture the extent of the experiments and the results. Those mentioned below are an incomplete set, and the present author assumes responsibility for any interpretive bias introduced by important omissions.

Transmission

OC-48 Sonet links were installed in four testbeds over distances of up to 2000 km, accelerating vendor development and carrier deployment of high speed Sonet equipment, establishing multiple-vendor Sonet interconnects, enabling discovery and resolution of standards implementation compatibility problems, and providing experience with Sonet error rates in an operational environment

A detailed study of striping over general ATM networks concluded that cell-based striping should be used which can be introduced at LAN-WAN connection points, in conjunction with destination host cell re-ordering and an ATM-layer synchronization scheme

Switching

Prototype high speed ATM switches were implemented and deployed for experiments in several of the testbeds, supporting 622 Mbps end-to-end switched links using both 155 Mbps striping and single-port 622 Mbps operation

A TDMA technique was developed and applied to tandem HIPPI switches to demonstrate packet-based quality of service (QoS) operation in HIPPI circuit-oriented switching environments, and a study of preemptive switching of variable length packets indicated a ten-fold reduction in processing requirements was possible relative to processor-based cell switching

Interworking

Three different designs were implemented to interwork HIPPI with wide area ATM networks over both Sonet and all-optical transmission infrastructure; explorations included the use of 4x155 Mbps striping and non-striped 622 Mbps access, local HIPPI termination

A HIPPI-Sonet gateway was implemented which allowed transfer of full 800 Mbps HIPPI rates across striped 155 Mbps wide area Sonet links; capabilities included variable bandwidth allocation of up to 1.2 Gbps and optional use of forward error correction, with a transfer rate of 790 Mbps obtained for HIPPI traffic

Host I/O

Several different testbed investigations demonstrated the feasibility of direct cell-based ATM host connections for workstation-class computers; this work established the basis for subsequent development of high speed ATM host interface chipsets by industry and provided an understanding of changes required to workstation I/O architectures for Gigabit networking.

TCP/IP investigations concluded that hardware checksumming and data-copying minimization were required by most testbed host architectures to realize transport rates of a few hundred Mbps or higher; full outboard protocol processing was explored for specialized host hardware architectures or as a workaround for existing software bottlenecks.

A 500 Mbps TCP/IP rate was achieved over a 1000-mile HIPPI/Sonet link using Cray supercomputers, and a 516 Mbps rate measured for UDP/IP workstation-based transport over ATM/Sonet; based on other workstation measurements it was concluded that, with a 4x processing power increase relative to the circa 1993 DEC Alpha processor used, a 622 Mbps TCP/IP rate could be achieved using internal host protocol processing and a hardware checksum while leaving 75% of the host processor available for application processing.

Measurements comparing the XTP transport protocol with TCP/IP were made using optimized software implementations on a vector Cray computer; the results showed TCP/IP provided greater throughput when no errors were present, but that XTP performed better at high error rates due to its use of a selective acknowledgment mechanism.

Presentation layer data conversions required by applications distributed over different supercomputers were found to be a major processing bottleneck; by exploiting vector processing capabilities, revisions to existing floating point conversion software resulted in a fifty-fold increase in peak transfer rates.

Experiments with commercial large-scale parallel processing architectures showed processor interconnection performance to be a major impediment to Gigabit I/O at the application level; an investigation of optimal data distribution strategies led to a selection of application control for data distribution within the processor array in conjunction with use of a reshuffling algorithm to remap the distribution for efficient I/O.

Work on distributed shared memory (DSM) for wide area Gigabit networks resulted in several latency-hiding strategies for dealing with large propagation delays, with relaxed cache synchronization resulting in significant performance improvements.

Network Management

In different QoS investigations, a real-time end-to-end protocol suite was developed and successfully demonstrated using video streams over HIPPI and other networks, and a ‘broker’ approach was developed for end-to-end/network QoS negotiations in conjunction with operating system scheduling for strict real-time constraints.

An evaluation of processing requirements for wide area QoS queuing in ATM switches, using a variation of the WFQ algorithm, found that a factor of 8 increase in processing speed was needed to achieve 622 Mbps port speeds relative to the i960/33 MHz processor used for the experiments.

A speedup mechanism was developed for lost packet recovery in high delay-bandwidth product networks using TCP's end-to-end packet window protocol.

An end-to-end time window approach using switch monitoring and feedback was developed and evaluated to provide high speed wide area network congestion control, and performed according to simulation-based predictions.

A control and monitoring subsystem was developed for real-time traffic measurement and characterization using carrier-based 622 Mbps ATM equipment, and was used to capture medical application traffic statistics; histograms showed the resulting ATM cell traffic to be highly bursty and contain oscillatory patterns.

A data generation and capture device for 800 Mbps HIPPI link traffic measurement and characterization was developed and commercialized, and was used for network debugging and traffic analysis; more generally, many network equipment problems were revealed through the use of real application traffic during testbed debugging phases.

Applications and Support Tools

Investigations using quantum chemical dynamics modeling, global climate modeling, and chemical process optimization modeling applications identified pipelining techniques and quantified speedup gains and network bandwidth requirements for distributed heterogeneous metacomputing using MIMD MPP, SIMD MPP, and vector machine architectures.

Most of the applications realized significant speedups, with a superlinear speedup of 3.3 achieved using two machines for the chemical dynamics application; other important benefits of distributed metacomputing such as large software program collaboration-at-a-distance were also demonstrated, and major advances made in understanding how to partition application software.

Several distributed applications involving human interaction in conjunction with large computational modeling were investigated. These included medical radiation therapy planning, exploration of large geophysical datasets, remote visualization of severe thunderstorm

modeling and other problems. The radiation therapy planning experiments successfully demonstrated the value of integrating high performance networking and computing for real-world applications; other interactive investigations similarly resulted in new levels of visualization capability, provided new techniques for distributed application communications and control, and provided important knowledge on problems which can prevent Gigabit speed operation.

A number of software tools were developed to support distributed application programming and execution in heterogeneous environments; these included systems for dynamic load balancing and checkpointing, program parallelization, communications and runtime control, collaborative visualization, and near-real-time data acquisition for progress monitoring and results analysis.

Legacy of the Gigabit Testbed Project

The Gigabit testbeds produced a demonstration of what could be done in a variety of application areas, and educated people in the research community, industrial sector, and government to provide a foundation for future phases of high performance communication and computing development.

UNIVERSITY RESEARCH EXAMPLES

This section will cover only two of many NSF-sponsored high performance networking projects underway in U.S. universities. The Washington University (Saint Louis, Missouri) Gigabit / ATM Test Kit project features research on gigabit networking and includes an outreach to other U.S. research institutions to join in on collaborative applied development of gigabit networking. The CAVE Automatic Virtual Environment which was developed by the Electronic Visualization Laboratory at the University of Illinois at Chicago has been replicated an increasing number of research projects. The CAVE depends on high data rates to render the visualization experience. Both of these activities present opportunities for international collaboration.

Washington University Gigabit/ATM Test Kit[2]

The gigabit ATM research program, now nearing completion at Washington University will conclude with the construction of a campus testbed network comprising several eight port ATM switches and compatible network interface cards for workstations and personal computers. These will be built around a set of four custom integrated circuits (three for the switch, one for the network interfaces) designed for this project. The gigabit network kits will include copies

[2]This description is excerpted from ;materials which can be found at http://www.arl.wustl.edu/~jst/gigatech/cfp.html, including other referenced reports.

of these switches and network interface cards. This section provides an overview of the Washington University ATM technology.

NSF has provided funding for the production of approximately 40 Gigabit Network Kits (see Figure 7, below) and to fund additional program activities associated with the distribution and use of the kits. The kits are intended to support a wide range of research agendas, including IP research, distributed systems, high performance computing, etc. To obtain a kit, eligible organizations will submit proposals describing their research objectives and explaining how the availability of one or more kits can help them achieve their goals. Proposals will be accepted from U.S. academic institutions. Proposals will be peer-reviewed, and the highest ranked proposals will be selected to receive kits. Program participants will also participate in additional activities, including attending a two week intensive course at the start of the program covering the gigabit network technology in detail, and subsequent workshops where they will report on results and share ideas with other program participants. The program will provide funding for two individuals from each participating organization to attend these events.

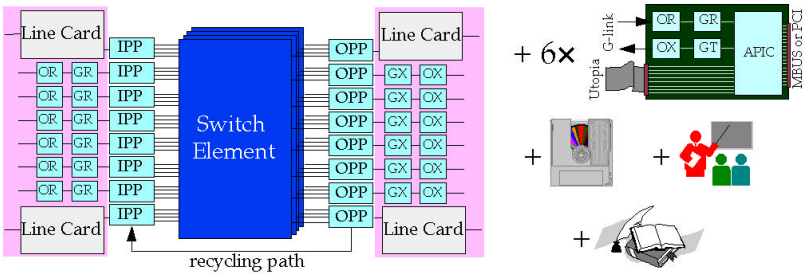

Figure 7. Washington University Gigabit Network Test Kit with the Washington University Gigabit Switch (WUGS) at left and ATM Port Interconnect Chip (APIC) at the right

Provisions are being made to allow other organizations (including, for example, non-U.S. universities) to purchase kits for research use (the current estimated cost is approximately $40,000). Professor Turner expressed a willingness to cooperate with efforts to internationalize the program. His suggested model for cooperative research would be for universities in other countries to propose companion programs like the Washington University program. Washington University would provide the hardware and software to them, but one university in each country should handle the coordination of activities (training sessions, local workshops, etc.) in that country. Specific arrangements for such purchases are still being developed. The deadline for U.S. proposals is the end of

January, 1997. Therefore, interested non-U.S. researchers should contact Professor Jonathan Turner[3] as soon as possible.

At the left side of the Figure 7 is a block diagram of the eight port ATM switch that will be included with the gigabit network kits. At the core of the switch is an eight port, shared buffer switch element made up of four chips, each of which switches an eight bit portion of the 32 bit wide internal data paths. Since the 32 bit data paths are clocked at a rate of 120 MHz, the raw data throughput is over 30 Gbps. Internal overhead in the switch reduces this to a usable data rate of 25.4 Gbps or about 3.2 Gbps per port, making the system fast enough for external links of up to 2.4 Gbps while operating at an occupancy of 75%. The switch element can accept cells with a pair of output addresses and is designed to forward copies of the cell to both specified outputs, or optionally, to all outputs in the range bounded by the specified outputs. The copied cells can (optionally) be recycled back to the corresponding input ports, allowing an arbitrary number of copies to be produced from a single input cell.

ATM port interface functions are provided by two additional chips, the Input Port Processor (IPP) and the Output Port Processor (OPP). The IPP includes a virtual path and virtual circuit translation table, used to determine the output port and new VPI/VCI for incoming cells. The OPP buffers cells awaiting transmission on the external links. The OPP and IPP are also linked by a direct data path which is used to allow cells to be recycled through the fabric multiple times to enable multicast transmission. This recycling data path, used in combination with the copy-by-two capability of the switch allows a multicast virtual circuit with f distinct outputs to be implemented using $\log_2 f$ passes through the switch.

The prototype switches being built at Washington University have six fixed transmission interfaces that operate at 1.2 Gbps and are based on Hewlett-Packard serial data link chips (called G-link chips). These chips provide a simple, inexpensive and higher performance alternative to Sonet OC-12C interfaces which have only recently become available. The prototype switches also have two ports that accept pluggable line cards, allowing some flexibility in interfacing to other types of network equipment. Washington University has designed a line card that supports a single OC-3C link and a line card that provides a 2.4 Gbps link using a pair of parallel fibers. Through a joint research project, Bellcore has also designed an OC-12C line card and an OC-48C line card should become available from Tektronix in the near future, through a similar joint project. Initial plans are to equip the kits with a pair of OC-3C line cards, but other options are under consideration.

The photographs in Figure 8 show the prototype switch (with the covers removed). The switches are packaged as ``deskside'' units, approximately 75 cm

[3]Personal communications with J. Turner. e-mail: jst@arl.wustl.edu (Jon Turner)

tall by 45 cm deep by 15 cm wide. In the left hand photo, the main system board, including the core ATM switching components and the six fixed transmission interfaces are visible. (In this photograph, only the switch element chips are present, although pin sockets for the IPP and OPP chips can be seen surrounding the switch elements.) The right hand photo shows the switch from the opposite side with two line cards plugged in. The line cards have purposely been made quite large, enabling the construction of experimental cards that incorporate added functionality (rate or credit-based flow control, IP routing features, video compression, data encryption, non-intrusive traffic monitoring, etc.).

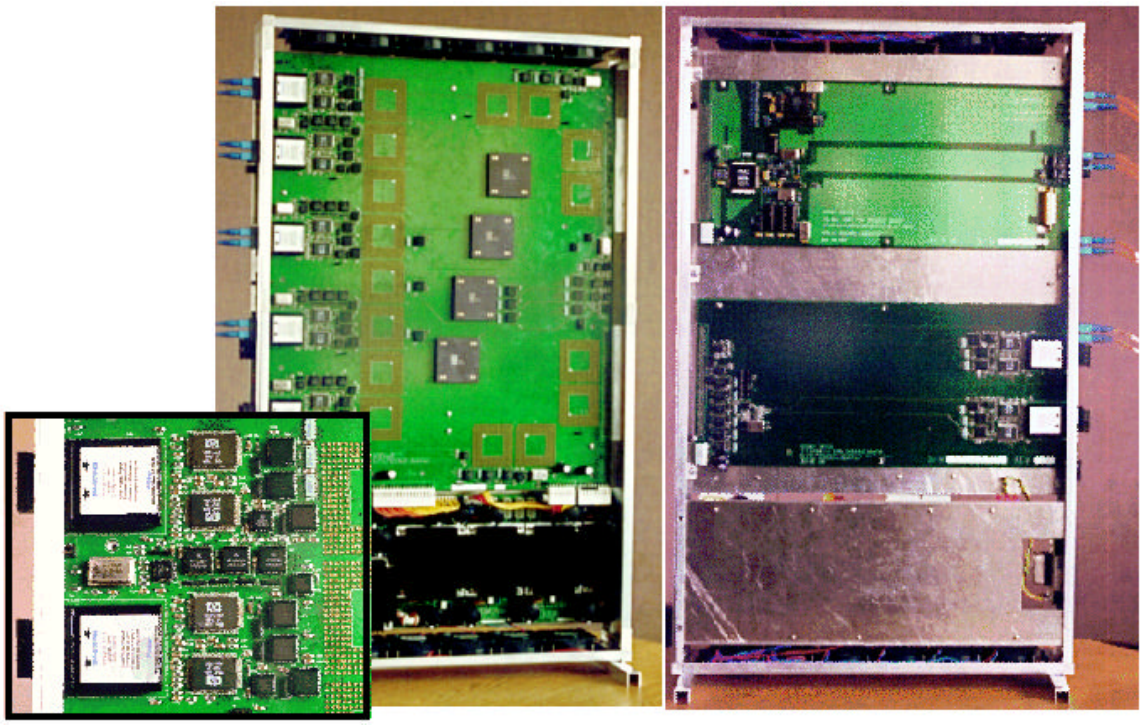

Figure 8. Washington University Prototype Gigabit Switch (covers removed)

The second key component of Washington University's gigabit ATM technology is a host-network interface device called the ATM Port Interconnect Chip (APIC), shown in Figure 9. The APIC has been designed with two ATM ports, allowing a series of APICs to be connected in a daisy-chain configuration, as shown in the upper part of the figure. This allows direct network-to-device communication, as well as device-to-device communication (that is, data need not pass through the computer system's main memory). This is particularly useful for applications like real-time video which need not pass through a general purpose processor on their way from network to display. Each of these ATM ports operates at up to 1.2 Gbps (full-duplex). and follows the Utopia interface specification which includes a simple back pressure flow control mechanism that can be used between APICs.

In addition, the APIC has an external memory interface port that is compatible with Sun's M-bus and with 32 bit and 64 bit PCI bus systems. The M-bus interface, allows the APIC to pass data directly to and from a Sun workstation's main memory, without being limited by the bandwidth of an intervening I/O bus controller, permitting peak data rates of over 1 Gbps.

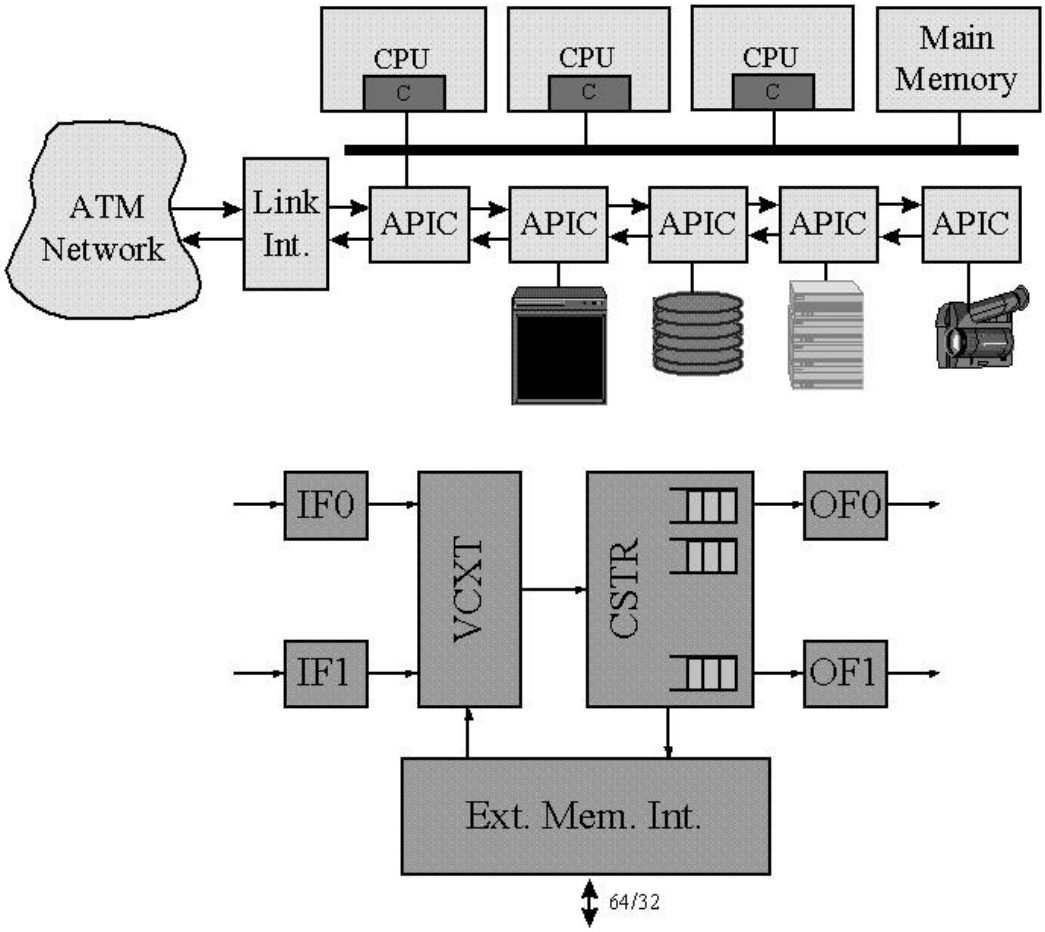

Figure 9. ATM Port Interconnect Chip (APIC)

The planned schedule for production and shipping of the gigabit network kits will make them available to users during the summer of 1997.

The APIC chip design is in the final stages and is expected to be completed and fabricated, before the end of 1996. In parallel with the APIC design, prototype boards have been built and tested to verify the state machines for the Mbus and PCI bus interfaces. The network interface cards for the APIC are currently being designed and will be completed well before the chips are delivered from the foundry.

The APIC software development is under way now, with the core components (device driver, system call interface, developer's library and raw ATM interface) scheduled to be completed and tested on an APIC simulator by the end of 1996 . The remaining components are to be completed during the first half of 1997.

University of Illinois CAVE Automatic Virtual Environment (CAVE)[4]

We now turn our attention from a networking tool kit to a computationally-intensive virtual reality (VR) environment construct that can easily consume large bandwidths in support of distributed applications. The CAVE (Figure 10) is a multi-person, room-sized, high-resolution, three-dimensional (3D) video and audio environment. In the current configuration, graphics are rear projected stereoscopically onto three walls and the floor, and viewed with stereo glasses. As a viewer wearing a location sensor moves within its display boundaries, the correct perspective and stereo projections of the environment are updated, and the image moves with and surrounds the viewer. The other viewers in the CAVE are like passengers in a bus, along for the ride!

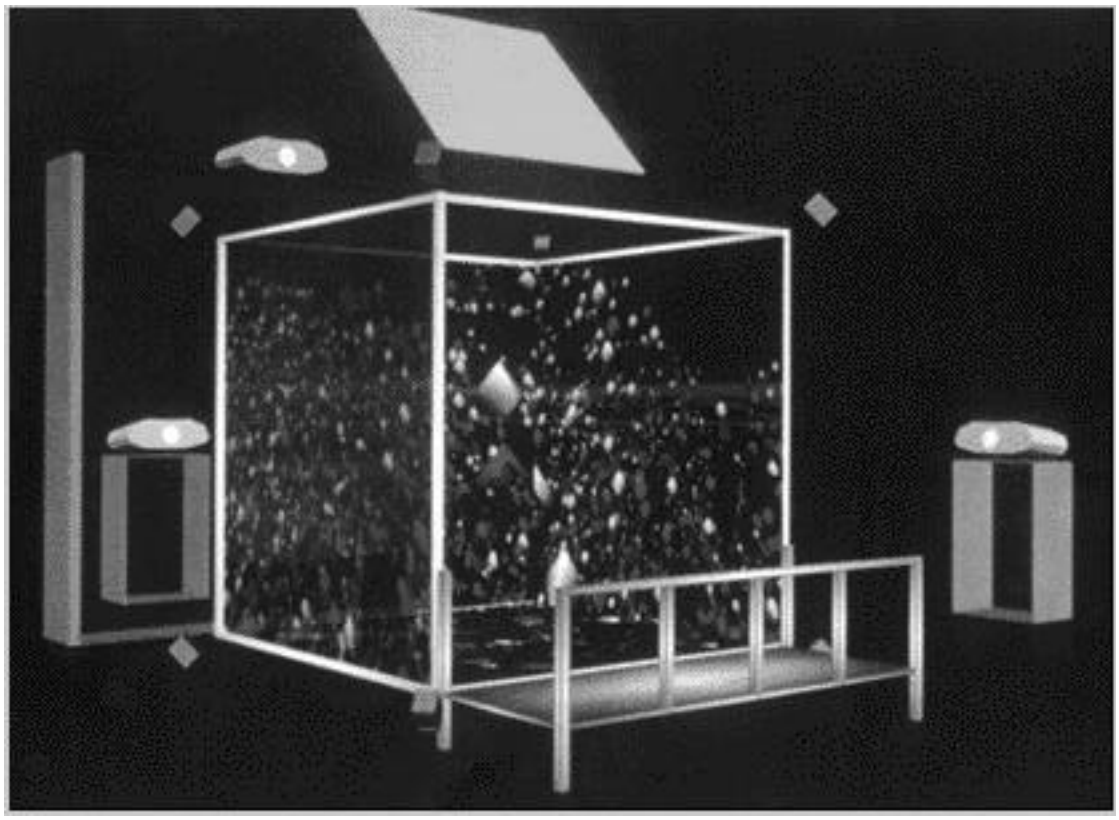

Figure 10. Depiction of a CAVE, showing the external rear-projection apparatus

More specifically, the CAVE is a theater 10 x10 x 9 feet, made up of two rear-projection screens for walls and a down-projection screen for the floor. Electrohome Marquis 8000 or 8500 projectors throw full-color workstation fields

[4]http://evlweb.eecs.uic.edu/EVL/VR/systems.html

(1024x768 stereo) at 96 Hz onto the screens, giving approximately 2,000 linear pixel resolution to the surrounding composite image. Computer-controlled audio provides a sonification capability to multiple speakers. A user's head and hand are tracked with Ascension tethered electromagnetic sensors. Stereographics' LCD stereo shutter glasses are used to separate the alternate fields going to the eyes. Two SGI Onyxes with Reality Engines or InfiniteReality Engines are used to create the imagery that is projected onto the walls and floor. The CAVE's theater area sits in a 30 x 20 x 13-foot light-tight room, provided that the projectors' optics are folded by mirrors.

One example of the use of the CAVE is transcontinental collaboration over high-speed and high-bandwidth networks connected to heterogeneous supercomputing resources and large data stores. Calvin (Collaborative Architecture Via Immersive Navigation) is a collaborative design laboratory which supports multiple participants in VR. Co-presence of designers is supported with Avatars, artificial graphical constructs which reproduce the actions of the co-participants in the virtual environment.[5] Another CAVE-to-CAVE application is a collaboration between Argonne National Laboratory and Nalco Fuel Tech to build an immersive interactive engineering tool for designing pollution control systems for commercial boilers and incinerators.[6]

The CAVE and its single-surface drafting-table-like cousin, the ImmersaDesk (ID), have been commercialized. They are sold by Pyramid Systems Inc.[7] For information on purchasing a CAVE or IDesk, or licensing CAVE software, contact Pyramid or Maggie Rawlings.[8]

NEW INTERNET INITIATIVES

In 1996, two initiatives for dramatically improving Internet performance and services were announced in the United States: a community of U.S. universities started on a path toward “Internet II,” and President Clinton launched the development of “Next Generation Internet.” While these two initiatives have some common objectives and elements, they are largely independent of each other. [9]

---

[5]Leigh, J. and Johnson, A.E., “Supporting Transcontnental Collaborative Work in Persistent Virtual Environments”, IEEE Computer Graphics and Applications, July 1996 pp 47-51; e-mail jleigh@eecs.uic.edu

[6]Diachin, D. *et al.,* “Remote Engineering Using CAVE-toCAVE Communications, note in IEEE Computer Graphics and Applications, July 1996, pg. 16 and http://www.mcs.anl.gov/home/freitag/nalco/nalco.html

[7]http://www2.pyramidsystems.com/psi/; e-mail:psi_sales@pyramidsystems.com

[8] e-mail: maggie@evl.eecs.uic.edu

[9]To round out North American developments, it should be noted that in Canada, the wheels are also in motion to create the "next generation" Internet network, dubbed “CA*net II.” See http://www.evert.com/canarie/nl9636sp.htm for details.

"Internet II" Pursued by University Community

Computing officials at 34 research universities have agreed to work together to build Internet II, a nationwide computing network that will operate at speeds several times faster than today's Internet. Internet II will be a pre-commercial, applications development project.[10] The universities have pledged to create a new organization, financed by membership fees, which will help to fund the effort. Partners in the effort will include computer and telecommunications companies, in addition to the higher education community. The new network will be used only for communications between the campuses that choose to participate – all messages heading toward other locations will still use the Internet. Educom VP Mike Roberts noted that the plan calls for involving as many corporate partners as possible: "The challenge here is not to invent it, but to integrate it." He adds that Internet II would solve the current "chicken-and-egg" problem, where the telecommunications companies are reluctant to invest in the next generation of network technology until there is a critical mass of potential users.[11]

As a precursor to the Internet II Project, The Monterey Futures Group, an outgrowth of several U.S.-based organizations concerned with the application of information technology to education, posited a comprehensive set of requirements for supporting higher education in the Year 2000.[12]

General Requirements:

The following general requirements were identified ... to meet higher education needs in the year 2000:

- Improved information security
- Authorization, and authentication capabilities
- Network management capabilities including performance audit

Quality of Service:

The quality of network service is becoming increasingly important as the network is relied upon for day-to-day and mission critical applications. Solutions and standards must be developed that insure adequate levels of:

- Latency and jitter
- Bandwidth interrogation capability and bandwidth reservation
- Guarantee of delivery

---

[10]See http://www.internet2.edu for more information about Internet II.

[11]Chronicle of Higher Education 11 October 96 A29

[12]"Toward a National Higher Education Networking Infrastructure", Part II of III: *White Paper on Telecommunications Requirements for a Virtual University*. Author: Doug Gale and many in the Monterey Futures Group (MFuG), URL:http://www.farnet.org/papers.html

Wide-area Communications:

...the Chicago Workshop [of the Monterey Futures Group] attendees developed the following estimates of the wide area connectivity requirements for higher education to meet its needs in the year 2000:

- Scaleable to serve all 3,600 institutions of higher education
- Supports integrated, real-time voice, video, and data
- OC-12 to OC-48 connections into the national backbone infrastructure for the research universities
- Supports individual streams at OC-3
- End-to-end Quality of Service (low latency and jitter)
- Supports multicasting
- Connections to other networks, both national and international

Campus Communications:

The following requirements represent the best estimates of the Chicago Workshop attendees of the campus-wide requirements to meet higher education needs in the year 2000:

- OC-12 (622 Mb/s) or greater bandwidths at the campus core
- OC-3 (155 Mb/s) or greater connecting the core to distributed locations (edge boxes)
- End-to-end Quality of Service
- Support hundreds to thousands of 1.5 Mb/s streams
- Multicast capabilities
- Symmetric connectivity

Community Infrastructure:

The following requirements represent the best estimates of the Chicago Workshop attendees of the community infrastructure requirements to meet higher education needs in the year 2000:

- Mb/s to 10 Mb/s into the home - Goal of connectivity in the home being the same as on campus
- Support 10BaseT connectivity
- Future target of Quality of Service

These requirements reflect estimates that could scale to meet the needs of all (about 3600) U.S. institutions of higher education. The Internet II Project is aimed at developing pre-commercial advanced capabilities for a much smaller community of perhaps 100 large, research-intensive universities in the U.S. and then transferring those capabilities to private sector broadband networking. Therefore, the Internet II Project itself would not necessarily implement gigabit networking. However, the eventual target community of all U.S. institutions of higher education is likely to require a multi-gigabit national backbone in order to meet its needs.

"Next Generation Internet" Launched by White House

The Administration's "Next Generation Internet" initiative has three goals:[13]

Connect universities and national labs with high-speed networks that are 100 - 1000 times faster than today's Internet

These networks will connect at least 100 universities and national labs at speeds that are 100 times faster than today's Internet, and a smaller number of institutions at speeds that are 1,000 times faster. ...

Promote experimentation with the next generation of networking technologies

For example, technologies are emerging that could dramatically increase the capabilities of the Internet to handle real-time services such as high quality video-conferencing. There are a variety of research challenges associated with increasing the number of Internet users by a factor of 100 that this initiative will help address. By serving as "testbeds", research networks can help accelerate the introduction of new commercial services.

Demonstrate new applications that meet important national goals and missions

Higher-speed, more advanced networks will enable a new generation of applications that support scientific research, national security, distance education, environmental monitoring, and health care. Below are just a few of the potential applications:

Health care: Doctors at university medical centers will use large archives of radiology images to identify the patterns and features associated with particular diseases. With remote access to supercomputers, they will also be able to improve the accuracy of mammographies by detecting subtle changes in three-dimensional images.

National Security: A top priority for the Defense Department is "dominant battlefield awareness," which will give the United States military a significant advantage in any armed conflict. This requires an ability to collect information from large numbers of high-resolution sensors, automatic processing of the data to support terrain and target recognition, and real-time distribution of that data to the warfighter. This will require orders of magnitude more bandwidth than is currently commercially available.

---

[13]"Background On Clinton-Gore Administration's Next-Generation Internet Initiative," and "Press Briefing By Mike McCurry, Education Secretary Dick Riley And Greg Simon, The Vice President's Chief Domestic Policy Advisor," The White House, Office of the Press Secretary, October 10, 1996

Distance Education: Universities are now experimenting with technologies such as two-way video to remote sites, VCR-like replay of past classes, modeling and simulation, collaborative environments, and online access to instructional software. Distance education will improve the ability of universities to serve working Americans who want new skills, but who cannot attend a class at a fixed time during the week.

Energy Research: Scientists and engineers across the country will be able to work with each other and access remote scientific facilities, as if they were in the same building. "Collaboratories" that combine video-conferencing, shared virtual work spaces, networked scientific facilities, and databases will increase the efficiency and effectiveness of our national research enterprise.

Biomedical Research: Researchers will be able to solve problems in large-scale DNA sequencing and gene identification that were previously impossible, opening the door to breakthroughs in curing human genetic diseases.

Environmental Monitoring: Researchers are constructing a "virtual world" to model the Chesapeake Bay ecosystem, which serves as a nursery area for many commercially important species.

Manufacturing engineering: Virtual reality and modeling and simulation can dramatically reduce the time required to develop new products.

Funding:

The Administration will fund this initiative by allocating $100 million for R&D and research networks to develop the Next Generation Internet, $70 million from the existing defense budget and $30 million from the domestic discretionary budget across a number of different programs. The $100 million will be funded under the high-performance computing program, and the funding will be spread across a number of agencies – the National Science Foundation, the Defense Advanced Research Projects Agency, the Department of Energy, NASA, and the National Institutes of Health. Other agencies may be involved in promoting specific applications related to their missions.

APPROACHING WIDE AREA GIGABIT TRANSMISSION RATES

Several recent developments illustrate that operational and demonstration telecommunications activities are within a factor of ten, and in some cases as little as two, of reaching gigabit speeds. U.S. examples are presented, but telecommunications in other countries have also progressed to the immediate sub-gigabit level.

National Backbones Ramp up the Sonet Hierarchy

The U.S. carrier, MCI, has completed its $60 million upgrade to its internetMCI backbone, effectively quadrupling speeds from 155 Mbps to 622 Mbps. Based on its estimate that Internet traffic is growing at about 30% per month, MCI reportedly plans to double its backbone capacity next year.[14] MCI also operates the very high speed Backbone Network Service (vBNS) in partnership with NSF. The current bandwidth of the vBNS is 155 Mbps, but MCI is in the process of upgrading the vBNS to 622 Mbps.

MCI has employed IP over ATM over Sonet to implement the service. However, the possibility of implementing IP directly over 155 Mbps Sonet has been demonstrated recently. On September 23, 1996 the Swedish carrier, Tele2 and the U.S. carrier, Sprint brought into operation the worlds first transatlantic 155 Mbps native IP service. The circuit runs between Sprint NAP in Pennsauken, New Jersey, USA and Tele2 in Sweden and uses Cisco packet-over-Sonet/SDH technology. Similarly, Sprint, plans to upgrade its Internet backbone from multiple DS3's to Sonet/OC3c in the next few months.

High Data Rate Satellite Experiments at 155 Mbps and Above

The Advanced Communication Technology Satellite

The Advanced Communication Technology Satellite (ACTS) shown in Figure 11 was developed by the National Aeronautics and Space Administration (NASA) and was launched in 1993. It operates in the Ka-band (30/20 GHz) where there is 2.5 GHz of spectrum available, and it employs the following on-orbit advanced technologies:

- very high-gain, multiple hopping beam antenna systems which permit smaller aperture Earth stations

- on-board baseband switching which permits interconnectivity between users at an individual circuit level

- a microwave switch matrix which enables gigabit/second communication between users.

[14] "Broadcasting & Cable", November 25, 1966, as quoted in e-mail list Edupage, 3 December 1996

ACTS has already demonstrated successful operation at modest data rates, conducting experiments in many fields such as banking, distance learning, telemedicine, and in military and mobile service. But, ACTS' most significant set of demonstrations, those at high data rate, are just getting underway.

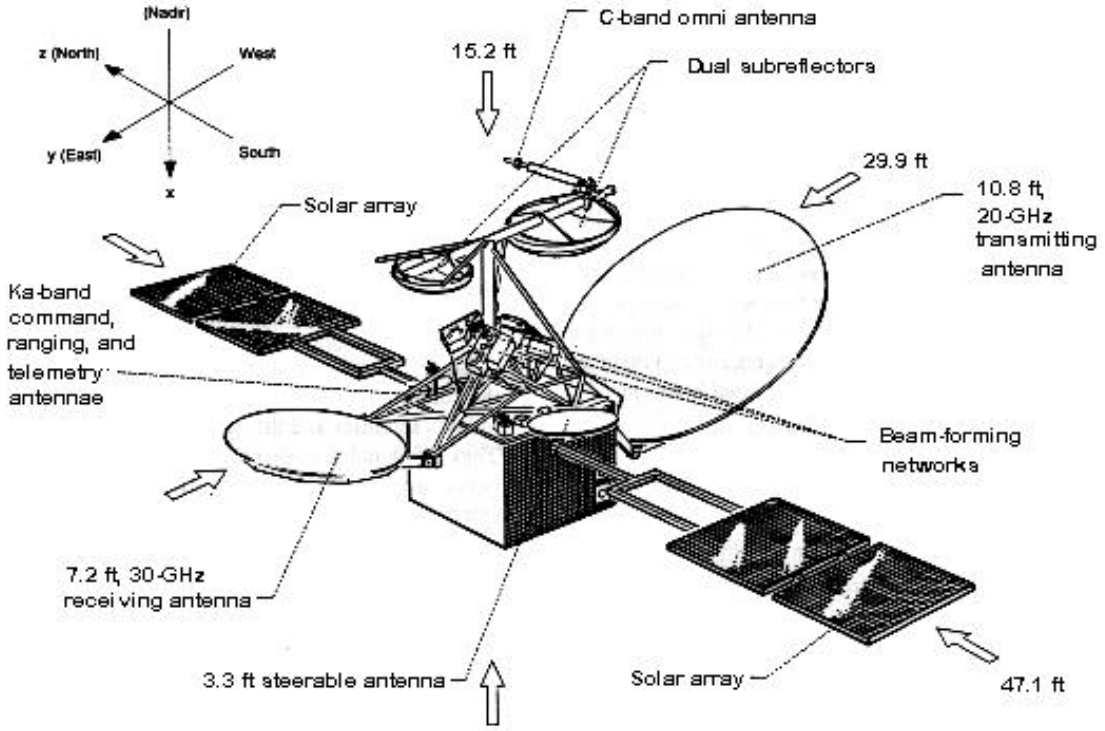

Figure 11. the ACTS Satellite

With five newly developed High Data Rate (HDR) terminals, ACTS is demonstrating the ability to transmit at up to OC-12 Sonet rates. One of ACTS' most demanding experiments is in supercomputing networking, in which a Cray supercomputer at the NASA Goddard Space Flight Center (near Washington, D.C.) is being connected with a Cray at the Jet Propulsion Laboratory (JPL, near Los Angeles, California) by HDR terminals at 155 Mbps. This experiment in global climate modeling will analyze and characterize the ocean-atmosphere interface. The oceanic information and scientists are at JPL, while the atmospheric information and scientists are at Goddard. In another experiment, now in progress, the Keck telescope in Hawaii is being connected to the astronomical data processing facility at the California Institute of Technology to perform a set of experiments in remote facility control and data visualization and analysis.

Soon, ACTS will assume an international role by serving as a link in a set of trans-Pacific HDR experiments originally planned as part of the Japan-U.S. Space Cooperative Program and now a part of the G7 Global Interoperability of Broadband Networking (GIBN) initiative. ACTS will connect the U.S. mainland and Hawaii, and Intelsat will connect Hawaii to Japan. The double-hop link will provide a very high definition video transmission test between Tokyo and Los Angeles for the Sony Corporation.[15]

ATM Research & Industrial Enterprise Study (ARIES) Project

Lest there be any doubt about the value of HDR satellite networking to industry, the ATM Research & Industrial Enterprise Study (ARIES) Project is offered as an example. The ARIES demonstration project is a collaboration between member companies of the American Petroleum Institute, NASA, the U.S. Department of Energy and the telecommunications industry.[16] ARIES has demonstrated both seismologic petroleum exploration and tele-medicine applications with an ocean-going vessel. The ACTS satellite is integrated seamlessly with terrestrial ATM networks to extend capabilities to platforms not served by terrestrial networks.

The rationale for providing a real-time satellite communications capability for petroleum industry operations is as follows:

- The petroleum industry does business where traditional service providers do not provide commercial telecommunications services

- During the exploration and production phase of operations, this involves generating very large data sets in very remote places, including offshore

- Information generated remotely needs to be delivered to computer centers and researchers at central locations

---

[15]Edelson, B.I., "High Data Rate Satellite Communications" prepared for presentation at the Second Ka-band Utilization Conference and SCGII Workshop, Florence, Italy, 24 September, 1996. Dr. Edelson can be contacted as edelson@seas.gwu.edu

[16]For information about ARIES, contact American Petroleum Institute ARIES Project Director
David R. Beering
Amoco Corporation
200 East Randolph Drive P0914D5,
Chicago, IL 60601 USA
e-mail: drbeering@amoco.com

- Every operational aspect of the petroleum industry deals with hazardous materials - if something goes wrong, it typically goes VERY wrong, VERY rapidly, [and broadband information should be made available to decision makers and experts quickly--thus the reason for demonstrating the handling of a simulated cardiac emergency aboard the seismologic collection ship].

- The capital cost of oil exploration in frontier exploration domains can reach into 10's of billions of dollars (U.S.)

- Since the cost of exploration in frontier exploration domains is so high, companies often form joint ventures to share the cost, risk , and eventual reward of the exploration endeavor

- In 1996 the oil industry spent approximately $60 Billion on oil exploration and production

- Of that, $2 Billion was spent collecting data offshore.[17]

Figures 12 and 13 illustrate seismic data collection from a ship at sea by means of ACTS. This has actually been demonstrated up to sea state 4 and data rates of several Mbps.

In 1997, ARIES will team with NASA to test the ACTS link at 622 Mbps between NASA's Lewis Research Center in Cleveland, Ohio, and JPL. The bandwidth-delay product at 622 Mbps (for a 520 ms round-trip delay) across the satellite is 40 Mbytes. The ARIES-LeRC-JPL team will use TCP/IP optimization schemes extensively to allow very fast data transfers over the satellite.

The test data will use very large chunks of the 3-Dimensional Synthetic Seismic Data Set (SSD). The sizes of the data sets are at least 120 MBytes. The team will also create several hybrid data sets in the 1-GByte-and-above range.[18]

[17]Beering, D.R. "ARIES," a PowerPoint Briefing, private communication

[18] e-mail message from drbeering@amoco.com (David Beering) "ARIES Year End 1996", 30 Dec 1966

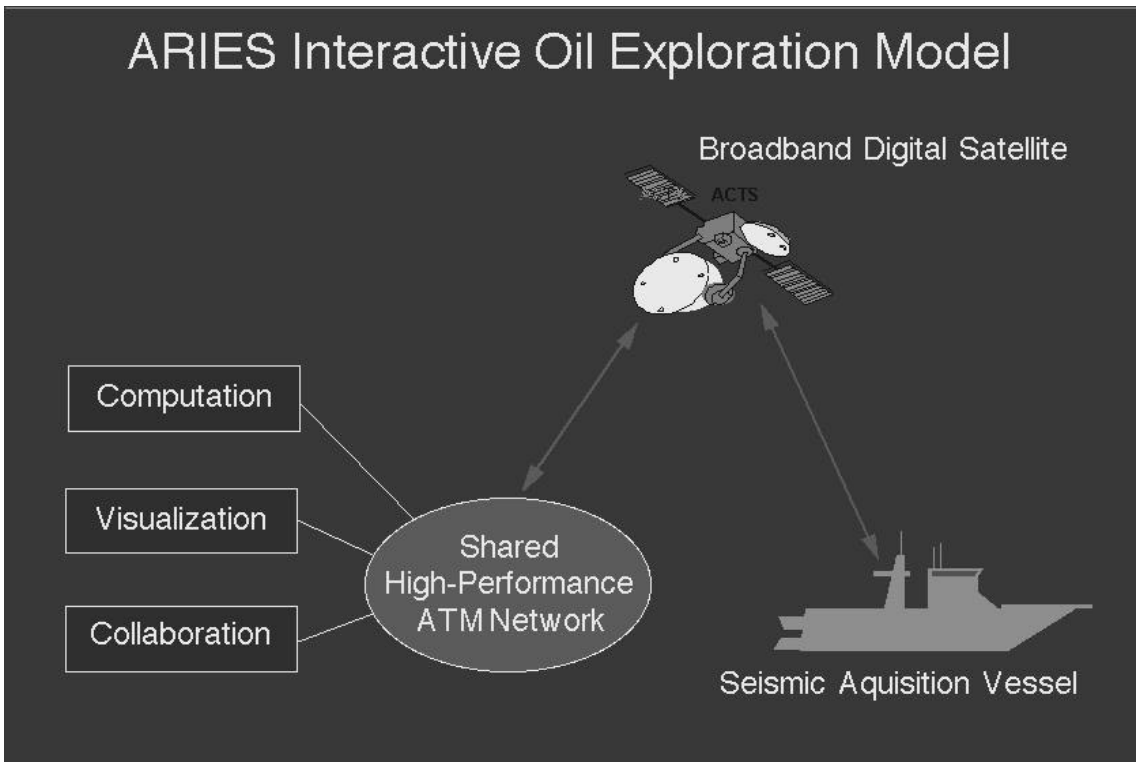

Figure 12. ARIES Interactive Oil Exploration Model

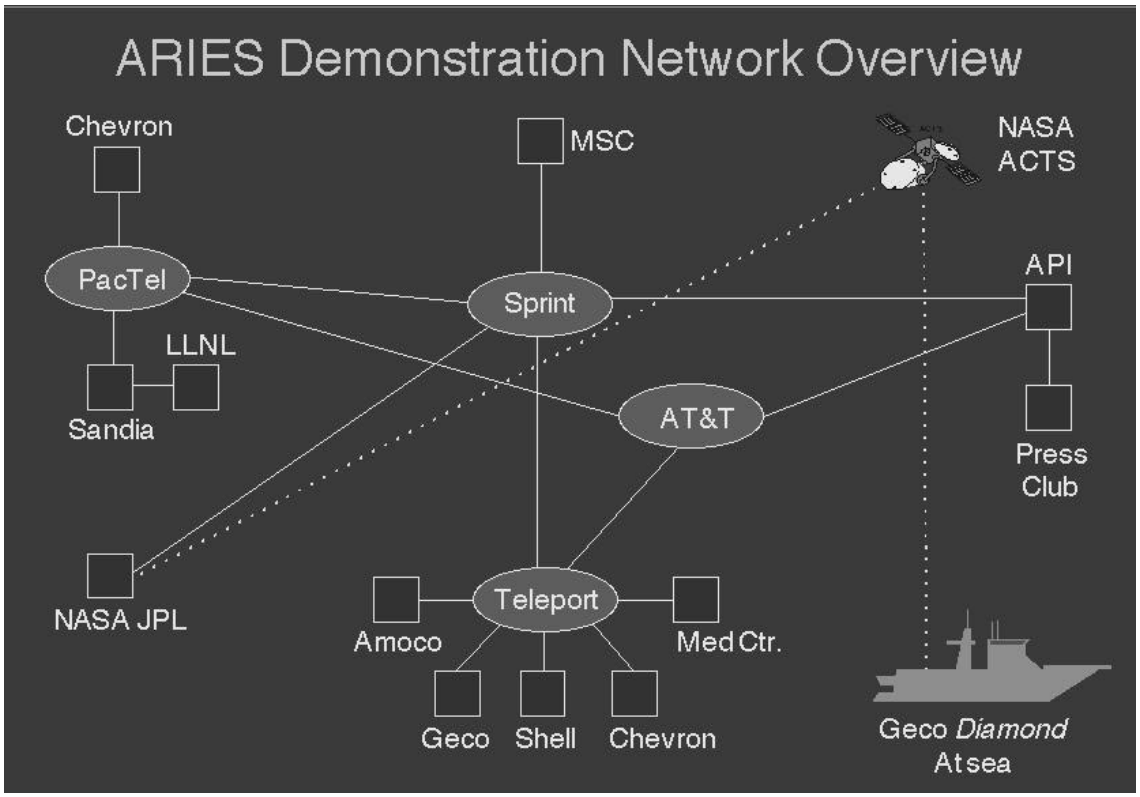

Figure 13. Details of the ARIES Demonstration Network

The technology developed and demonstrated in the ACTS Program provides a base for future satellite development in the gigabit regime. At least one commercial venture (Teledesic) plans to offer gigabit data rates by employing inter-satellite links at Low Earth Orbital altitudes. Japan is engaged in the development of a multi-Gbps data communications satellite. So, there is little doubt that the gigabit barrier for satellite-based networking will be challenged, if not broken, in the next several years.

Internet Routing and Switching Rises to the Challenge

Among the factors that threaten the future operation of the Internet are circuit capacity, address space depletion, and routing overload, that is, the [in-]ability of routers to keep up with both the huge amount of routing information needed to keep routes updated and also to forward customer data. Circuit capacity in many cases, and certainly in the case of backbones in industrialized countries, is likely to be a matter of economics at least for the next several years; that is, the capacity is likely to be there if Internet service providers (ISPs) can pay for it. Current address conservation measures with the present version of the Internet protocol (IPv4) seem to have slowed the address depletion sufficiently until the new version of the Internet Protocol (IPv6) with much larger address space can be phased into production. However, the outlook has not been so promising for "central office" class Internet routers that will be needed to support the backbones of large ISPs. A challenge facing the Internet routing fabric is the ability of the routers to be able to route datagrams while keeping up with the dynamic routing updates as the routing data base continues to grow.

Terabit Per Second Router Example

It has been said that the dominant router vendors, while aware of the inadequacies of router performance for major backbones, had more incentive to concentrate on enterprise routers whose market potential is in the hundreds of thousands of units than on central office routers whose market potential may number only in the thousands of units. The dominant router vendors are said to be pursuing the high-end central office class router development now, though their sense of priority is not known to the author. However, several new entrepreneurial companies are also said to be pursuing the high-end router product with all the vigor of startup businesses. One such company that has declared its intent openly is Pluris, Inc., "The Terabit Routing Company." [19]

---

[19] Pluris, Inc.
260 Sheridan Ave, Suite 212
Palo Alto, CA 94306
+1 (415) 328-8000
+1 (415) 654-9222 FAX
e-mail info@pluris.com
http://www.pluris.com

Vadim Antonov, President of Pluris, makes the case for the necessity of routers at terabit per second (Tbps) speeds.[20] He offers the graph reproduced in Figure 14. It shows that major ISP backbone speeds (squares) are likely to overtake the line speeds supported by routers (circles) sometime in 1997. His claim is that the new Pluris router will be capable of supporting the fastest backbone speeds well into the next decade. A description of the Pluris Massively Parallel Router (MPR), condensed from Antonov's paper, is presented below.[21]

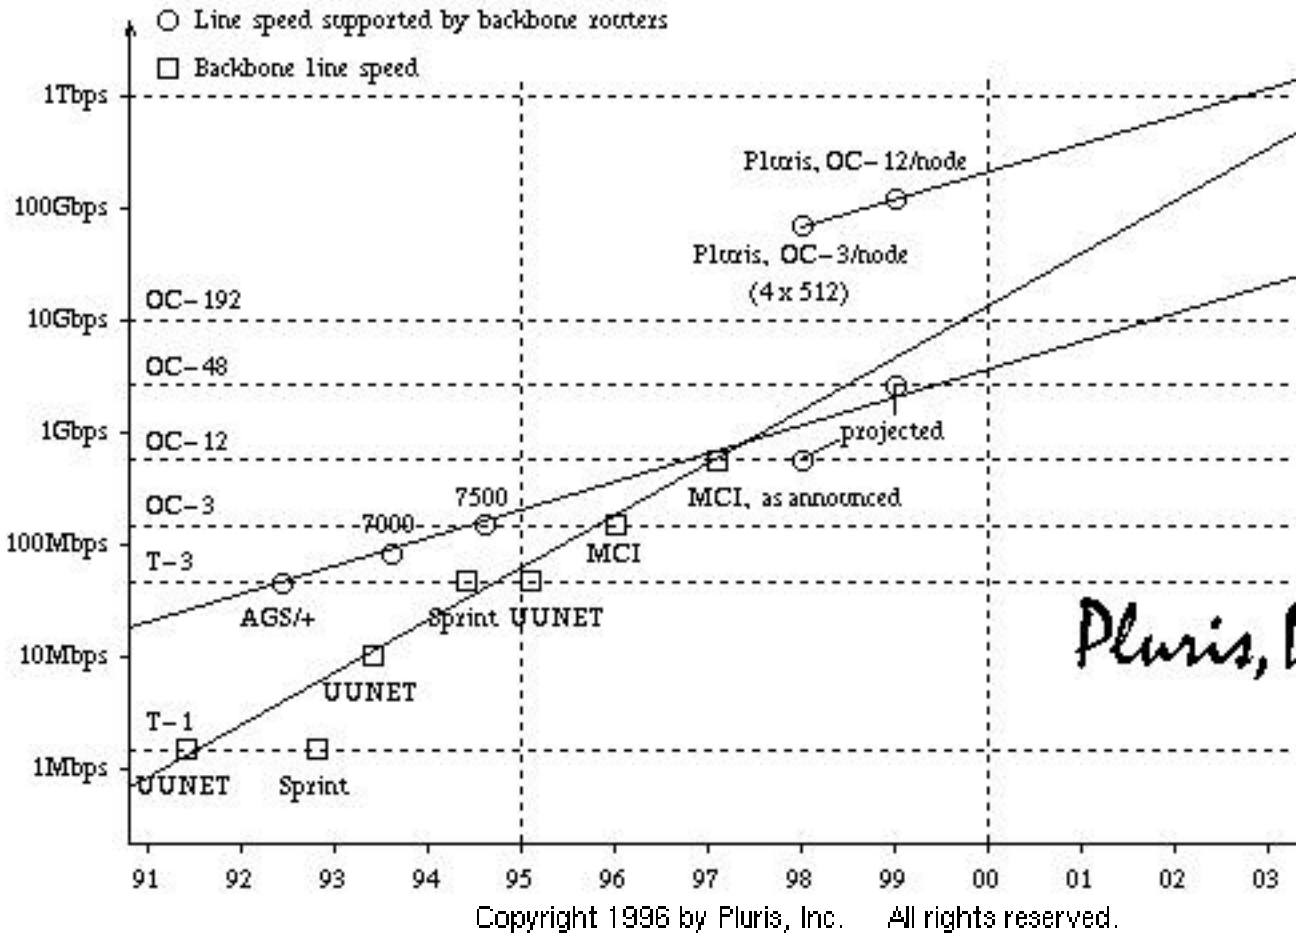

Figure 14. Internet backbone line speed and conventional router port speeds

The Pluris approach is based on Antonov's observation that although aggregate data rates of Internet traffic are skyrocketing, the bandwidth of individual communication sessions remains relatively small (in fact, it cannot grow faster than the performance of host computers). This means that a high aggregate

[20]http://www.pluris.com/terabit.html

[21]Pluris Massively Parallel Routing (White Paper)" http://www.pluris.com/wp/index.html

routing capacity can be achieved by distributing the paths of packets in those connections between a large number of medium-performance routing engines.

Overview of Pluris Router

A Pluris MPR is composed of a large number of such routing engines (which Pluris calls "processing nodes") Figure 15,[22] communicating with each other via a linearly scalable high-speed data interconnect. Such interconnects are relatively well-understood technology which is commonly used in loosely-coupled massively parallel computers.

The processing nodes are connected via low-speed lines to a number of synchronous multiplexers which combine the low-speed data streams into high-speed streams on backbone. Every processing node has its own copy of the forwarding table (that table, compared to BGP routing information bases, is not large).

Instead of conventional single-step IP routing (i.e. determination of exit interface from destination address), Pluris' process performs two steps for each packet: at the first step the exit high-speed communication line is determined, and at the second step one of the low-speed lines corresponding to the exit high-speed line is selected. The packet is then sent through the data interconnect to the processing node corresponding to the selected low-speed communication line.

The selection is made by computing a hash function from the packet's source and destination addresses and, optionally, port numbers. The use of the hash function from values of packet's fields which are invariant for all packets within a single TCP (or any other transport protocol) session guarantees that all those packets will follow the same path, and therefore will not be reordered. Pluris claims that hashing effectively randomizes packet routes, so the load is uniformly distributed between all participating processing nodes and low-speed lines.

This, together with linear scalability of the data interconnect, means that the aggregate capacity of the massively parallel packet router can be increased nearly indefinitely by simple addition of processing nodes. The only high-speed circuitry is in the synchronous multiplexers, and Pluris claims that circuitry is much simpler and cheaper than hardware implementations of IP routing or ATM switching. In fact, since Pluris routers treat high-speed backbone links as quantities of parallel low-speed circuits, a number of parallel multiplexed high-speed lines (for example, different strands of fiber in a cable, or different wavelength channels) can be combined into a single very high-speed communication line. Thus, the capacity of a network built using massively parallel routers is not limited by capacity of any physical component.

---

[22]Figures 15, 16, and 17, *op. cit.*

The Pluris MPR is a collection of single-board computers (processing nodes) and a proprietary data interconnect. Each (first generation) processing node has 16 or more megabytes of dynamic random access memory (DRAM) and a 100+ MHz general-purpose microprocessor, sufficient to route IP packets at OC-3c speed (155 Mbps). The second generation of processing nodes will support at least OC-12 per node, and will be compatible with the first generation nodes (Figure 15):

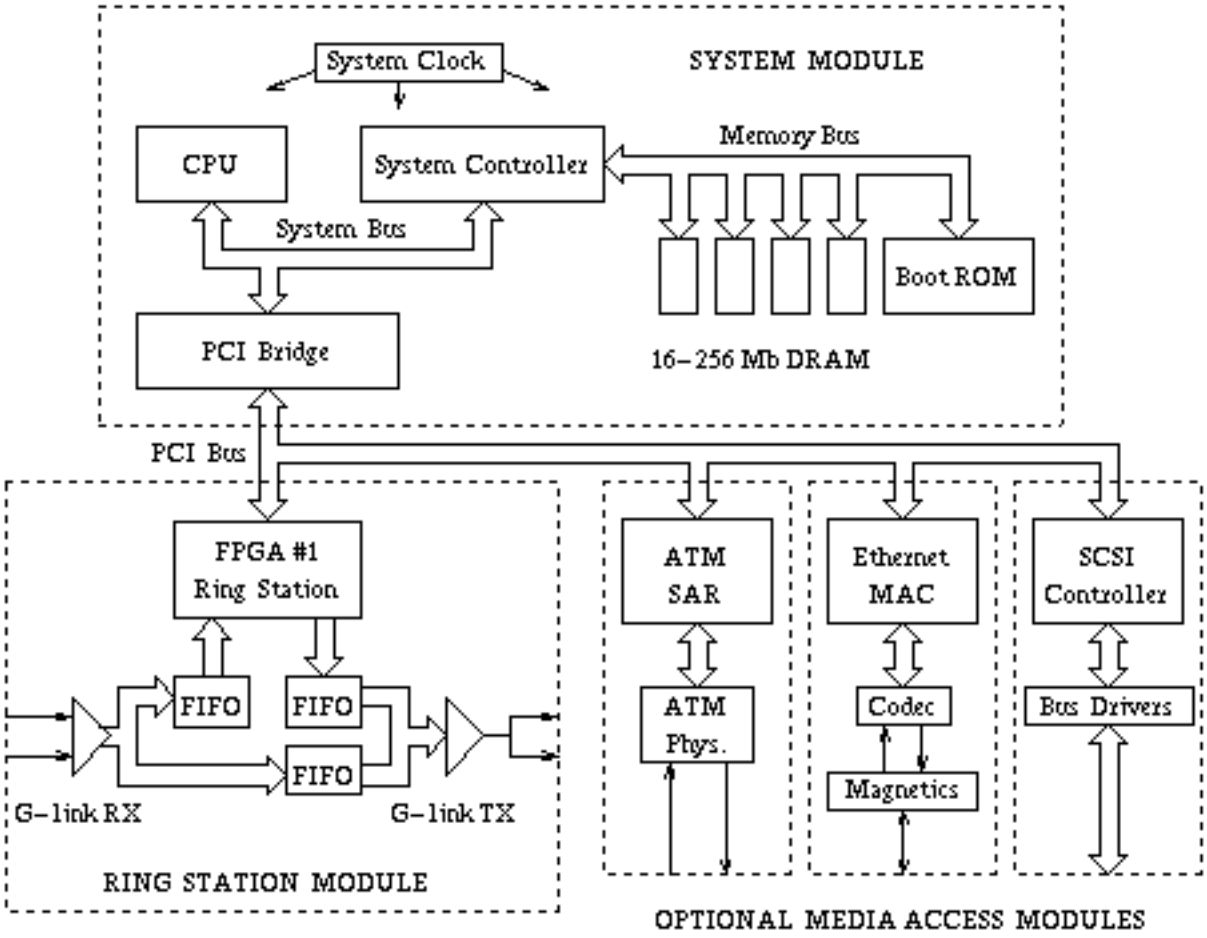

Figure 15. Block Diagram of the Pluris Central Processing Node

The only unusual feature of a processing node is the ring station module, connecting the processing node to the high-speed data interconnect.

The data interconnect is a patent-pending Self-Healing Butterfly Switch based on 1.2 Gbps serial communication lines (Figure 16):

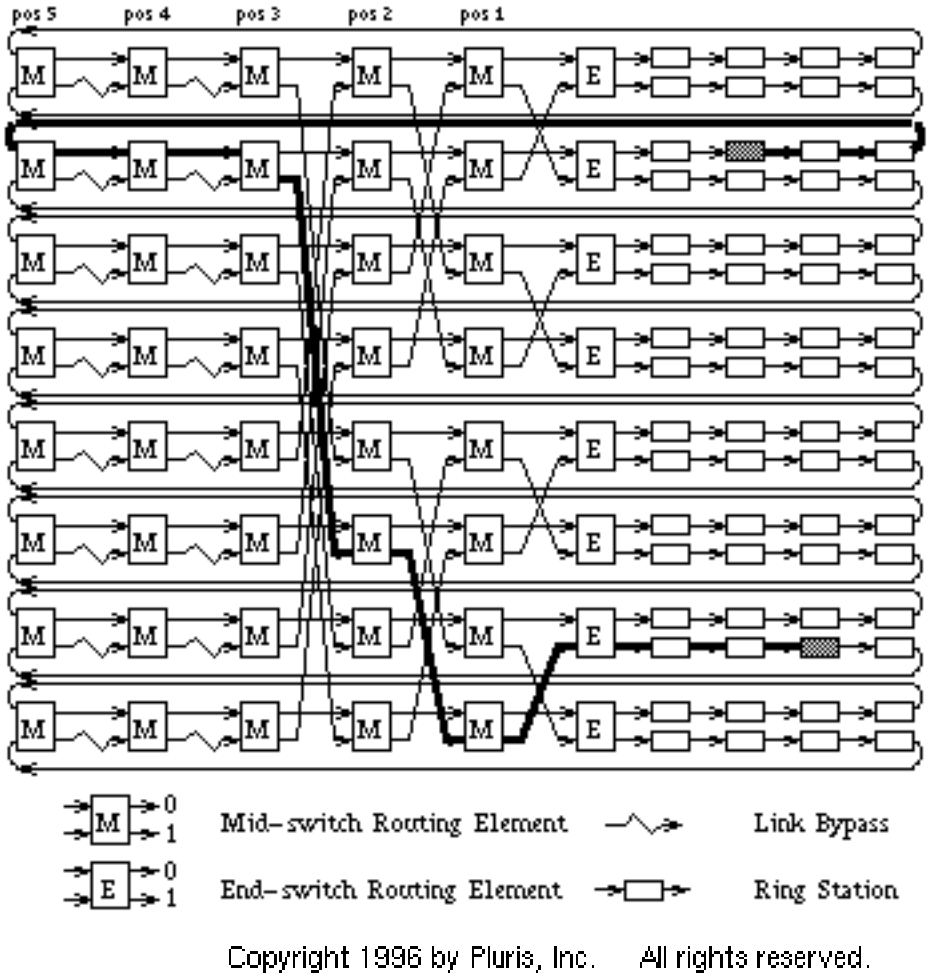

Figure 16. Self-Healing Butterfly Switch Design of the Router

The Pluris butterfly switch is fault-tolerant (the diagram above does not show secondary links), so the packets are automatically rerouted in case of failures in links or routing elements. Usage of radio-frequency serial lines reduces the amount of wiring between card cages. Every card cage has 16 processing nodes, and switch circuitry in additional inter shelf link boards. All boards are hot-swappable; also every card cage has redundant power supplies.

When several card cages are interconnected to form a larger system, the wiring is similar to the wiring of a hypercube-based massively parallel computer (Figure 17):

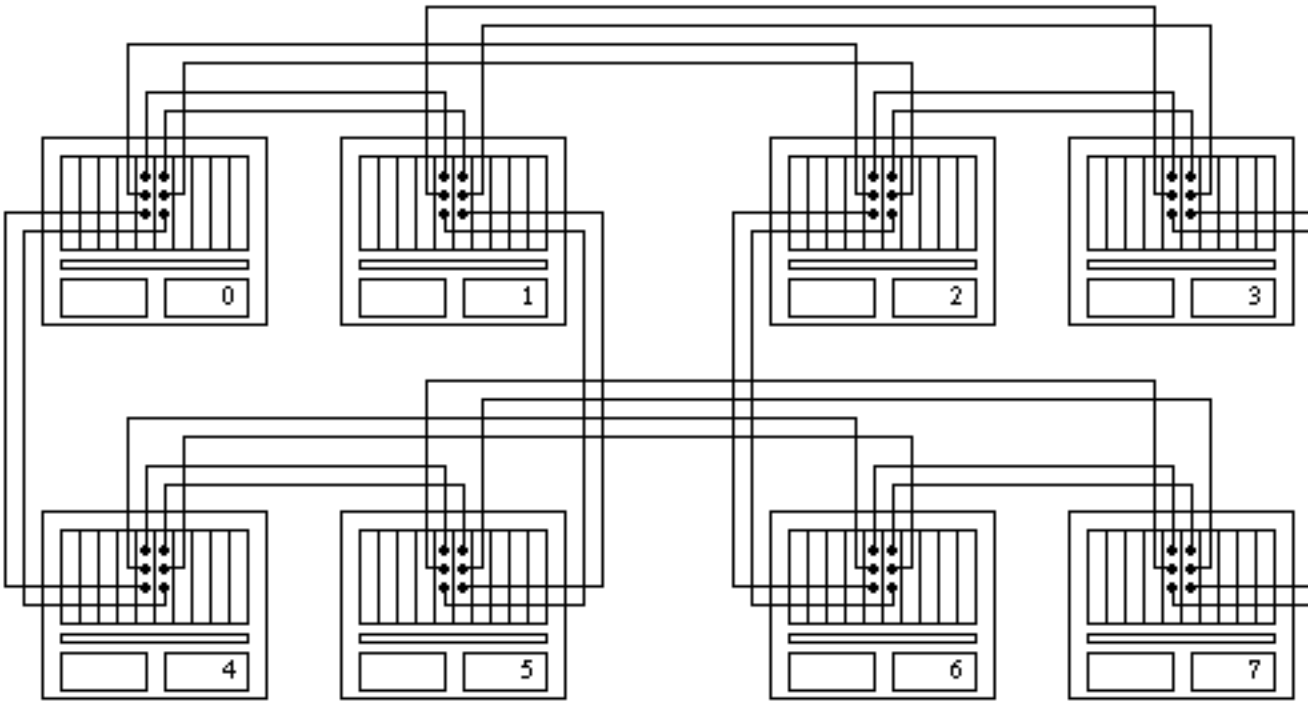

Figure 17. Wiring Schematic of a 128-node Parallel Processing Router

One or several dedicated processing nodes equipped with 64-256 Mb of DRAM are used for performing routing protocols. When several such nodes are used, the output of every protocol engine is broadcast to all forwarding nodes, so if a protocol engine node fails, or is removed, the operation will continue. A failure of a forwarding node only causes reduction of throughput, but not interruption of service.

Performance Claims

Pluris claims that although its MPR is a high-performance machine, it is entirely composed of off-the-shelf integrated circuits, making it a comparatively low-cost and reliable device. The maximum capacity of the MPR is limited solely by the maximum length of coaxial cables interconnecting parts of the machine. The present design is capable of housing 16 thousand processing nodes in 64 open racks arranged in 4 rows, to achieve the aggregate routing capacity of 2.4 Tbps (or 7 billion packets per second).

CONCLUSIONS

The author has borrowed liberally from the work of colleagues engaged in university-and-industry-partnered gigabit networking research, commercial product development, advanced network demonstration projects and satellite-based broadband technology. The examples offered illustrate the theme that wide area gigabit networking services and applications that depend on those services are close to being realized in the U.S. Opportunities for collaboration with international partners exist and are already underway in several areas.

AUTHOR INFORMATION

Steven Goldstein[23] is the U.S. National Science Foundation's program director for international networking. He has managed the development of international connectivity on behalf of the U.S. Research and Education community for the last seven years, during which time scores of countries have connected to the "International Connections Management for NSFNET" (ICM) infrastructure, and bandwidths of transatlantic links have grown from 64 kbps to 34 Mbps (1995) and 155 Mbps (1996). He also chairs the U.S. Delegation to the G7 Global Information Society Project entitled "Global Interoperability for Broadband Networking (GIBN)." He holds S.B. And S.M. Degrees in Physics from the Massachusetts Institute of Technology and the Ph.D. in Engineering and Public Policy from Carnegie-Mellon University.

[23]Steven N. Goldstein
National Science Foundation
4201 Wilson Boulevard, Room 1175
Arlington, Virginia 22230 USA
+1 703 306 1949 voice
+1 703 306 0621 FAX
Internet: sgoldste@nsf.gov